arXiv:1601.02387v1 [stat.CO] 11 Jan 2016

Bounding errors of Expectation-Propagation

Guillaume Dehaene
University of Geneva
guillaume.dehaene@gmail.com

Simon Barthelmé
CNRS, Gipsa-lab
simon.barthelme@gipsa-lab.fr

Abstract

Expectation Propagation is a very popular algorithm for variational inference, but comes with few theoretical guarantees. In this article, we prove that the approximation errors made by EP can be bounded. Our bounds have an asymptotic interpretation in the number n of datapoints, which allows us to study EP's convergence with respect to the true posterior. In particular, we show that EP converges at a rate of $\mathcal{O}(n^{-2})$ for the mean, up to an order of magnitude faster than the traditional Gaussian approximation at the mode. We also give similar asymptotic expansions for moments of order 2 to 4, as well as excess Kullback-Leibler cost (defined as the additional KL cost incurred by using EP rather than the ideal Gaussian approximation). All these expansions highlight the superior convergence properties of EP. Our approach for deriving those results is likely applicable to many similar approximate inference methods. In addition, we introduce bounds on the moments of log-concave distributions that may be of independent interest.

Introduction

Expectation Propagation (EP, 1) is an efficient approximate inference algorithm that is known to give good approximations, to the point of being almost exact in certain applications [2, 3]. It is surprising that, while the method is empirically very successful, there are few theoretical guarantees on its behavior. Indeed, most work on EP has focused on efficiently implementing the method in various settings. Theoretical work on EP mostly represents new justifications of the method which, while they offer intuitive insight, do not give mathematical proofs that the method behaves as expected. One recent breakthrough is due to Dehaene and Barthelmé [4] who prove that, in the large data-limit, the EP iteration behaves like a Newton search and its approximation is asymptotically exact. However, it remains unclear how good we can expect the approximation to be when we have only finite data. In this article, we offer a characterization of the quality of the EP approximation in terms of the worst-case distance between the true and approximate mean and variance.

When approximating a probability distribution $p(x)$ that is, for some reason, close to being Gaussian, a natural approximation to use is the Gaussian with mean equal to the mode (or argmax) of $p(x)$ and with variance the inverse log-Hessian at the mode. We call it the Canonical Gaussian Approximation (CGA), and its use is usually justified by appealing to the Bernstein-von Mises theorem, which shows that, in the limit of a large amount of independent observations, posterior distributions tend towards their CGA. This powerful justification, and the ease with which the CGA is computed (finding the mode can be done using Newton methods) makes it a good reference point for any method like EP which aims to offer a better Gaussian approximation at a higher computational cost. In section 1, we introduce the CGA and the EP approximation. In section 2, we give our theoretical results bounding the quality of EP approximations.

1 Background

In this section, we present the CGA and give a short introduction to the EP algorithm. In-depth descriptions of EP can be found in Minka [5], Seeger [6], Bishop [7], Raymond et al. [8].

1.1 The Canonical Gaussian Approximation

What we call here the CGA is perhaps the most common approximate inference method in the machine learning cookbook. It is often called the "Laplace approximation", but this is a misnomer: the Laplace approximation refers to approximating the integral $\int p$ from the integral of the CGA. The reason the CGA is so often used is its compelling simplicity: given a target distribution $p(x) = \exp(-\phi(x))$, we find the mode $x^\star$ and compute the second derivatives of ϕ at $x^\star$:

$$\begin{aligned} x^\star &= \text{argmin}\phi(x) \\ \beta^\star &= \phi''(x^\star) \end{aligned}$$

to form a Gaussian approximation $q(x) = \mathcal{N}\left(x|x^\star, \frac{1}{\beta^\star}\right) \approx p(x)$. The CGA is effectively just a second-order Taylor expansion, and its use is justified by the Bernstein-von Mises theorem [9], which essentially says that the CGA becomes exact in the large-data (large-n) asymptotic limit. Roughly, if $p_n(x) \propto \prod_{i=1}^n p(y_i|x)\, p_0(x)$, where $y_1 \dots y_n$ represent independent datapoints, then $\lim_{n\to\infty} p_n(x) = \mathcal{N}\left(x|x_n^\star, \frac{1}{\beta_n^\star}\right)$ in total variation.

1.2 CGA vs Gaussian EP

Gaussian EP, as its name indicates, provides an alternative way of computing a Gaussian approximation to a target distribution. There is broad overlap between the problems where EP can be applied and the problems where the CGA can be used, with EP coming at a higher cost. Our contribution is to show formally that the higher computational cost for EP may well be worth bearing, as EP approximations can outperform CGAs by an order of magnitude. To be specific, we focus on the *moment estimates* (mean and covariance) computed by EP and CGA, and derive bounds on their distance to the true mean and variance of the target distribution. Our bounds have an asymptotic interpretation, and under that interpretation we show for example that the mean returned by EP is within an order of $\mathcal{O}\left(n^{-2}\right)$ of the true mean, where n is the number of datapoints. For the CGA, which uses the mode as an estimate of the mean, we exhibit a $\mathcal{O}\left(n^{-1}\right)$ upper bound, and we compute the error term responsible for this $\mathcal{O}\left(n^{-1}\right)$ behavior. This enables us to show that, in the situations in which this error is indeed $\mathcal{O}\left(n^{-1}\right)$, EP is better than the CGA.

1.3 The EP algorithm

We consider the task of approximating a probability distribution over a random-variable $\mathcal{X} : p(x)$, which we call the *target distribution*. $\mathcal{X}$ can be high-dimensional, but for simplicity, we focus on the one-dimensional case. One important hypothesis that makes EP feasible is that $p(x)$ factorizes into n *simple* factor terms:

$$p(x) = \prod_i f_i(x)$$

EP proposes to approximate each $f_i(x)$ (usually referred to as *sites*) by a Gaussian function $q_i(x)$ (referred to as the *site-approximations*). It is convenient to use the parametrization of Gaussians in terms of natural parameters:

$$q_i(x|r_i, \beta_i) \propto \exp\left(r_i x - \beta_i \frac{x^2}{2}\right)$$

which makes some of the further computations easier to understand. Note that EP could also be used with other exponential approximating families. These Gaussian approximations are computed iteratively. Starting from a current approximation $(q_i^t(x|r_i^t, \beta_i^t))$, we select a site for update with index i. We then:

- Compute the *cavity* distribution $q^t_{-i}(x) \propto \prod_{j \neq 1} q^t_j(x)$. This is very easy in natural parameters:

$$q_{-i}(x) \propto \exp\left(\left(\sum_{j \neq i} r^t_j\right) x - \left(\sum_{j \neq i} \beta^t_j\right) \frac{x^2}{2}\right)$$

- Compute the *hybrid* distribution $h^t_i(x) \propto q^t_{-i}(x) f_i(x)$ and its mean and variance
- Compute the Gaussian which minimizes the Kullback-Leibler divergence to the hybrid, ie the Gaussian with same mean and variance:

$$\mathcal{P}(h^t_i) = \underset{q}{\text{argmin}} \left(KL\left(h^t_i | q\right)\right)$$

- Finally, update the approximation of f_i:

$$q^{t+1}_i = \frac{\mathcal{P}(h^t_i)}{q^t_{-i}}$$

  where the division is simply computed as a subtraction between natural parameters

We iterate these operations until a fixed point is reached, at which point we return a Gaussian approximation of $p(x) \approx \prod q_i(x)$.

1.4 The "EP-approximation"

In this work, we will characterize the quality of an EP approximation of $p(x)$. We define this to be any fixed point of the iteration presented in section 1.3, which could all be returned by the algorithm. It is known that EP will have at least one fixed-point [1], but it is unknown under which conditions the fixed-point is unique. We conjecture that, when all sites are log-concave (one of our hypotheses to control the behavior of EP), it is in fact unique but we can't offer a proof yet. If $p(x)$ isn't log-concave, it is straightforward to construct examples in which EP has multiple fixed-points. These open questions won't matter for our result because we will show that all fixed-points of EP (should there be more than one) produce a good approximation of $p(x)$.

Fixed points of EP have a very interesting characterization. If we note q^*_i the site-approximations at a given fixed-point, h^*_i the corresponding hybrid distributions, and q^* the global approximation of $p(x)$, then the mean and variance of all the hybrids and q^* is the same[1]. As we will show in section 2.2, this leads to a very tight bound on the possible positions of these fixed-points.

1.5 Notation

We will use repeatedly the following notation. $p(x) = \prod_i f_i(x)$ is the target distribution we want to approximate. The sites $f_i(x)$ are each approximated by a Gaussian site-approximation $q_i(x)$ yielding an approximation to $p(x) \approx q(x) = \prod_i q_i(x)$. The hybrids $h_i(x)$ interpolate between $q(x)$ and $p(x)$ by replacing one site approximation $q_i(x)$ with the true site $f_i(x)$.

Our results make heavy use of the log-functions of the sites and the target distribution. We note $\phi_i(x) = -\log\left(f_i(x)\right)$ and $\phi_p(x) = -\log\left(p(x)\right) = \sum \phi_i(x)$. We will introduce in section 2 hypotheses on these functions. Parameter β_m controls their minimum curvature and parameters K_d control the maximum d^{th} derivative.

We will always consider fixed-points of EP, where the mean and variance under all hybrids and $q(x)$ is identical. We will note these common values: μ_{EP} and v_{EP}. We will also refer to the third and fourth centered moment of the hybrids, denoted by m^i_3, m^i_4 and to the fourth moment of $q(x)$ which is simply $3v^2_{EP}$. We will show how all these moments are related to the true moments of the target distribution which we will note μ, v for the mean and variance, and m^p_3, m^p_4 for the third and fourth moment. We also investigate the quality of the CGA: $\mu \approx x^\star$ and $v \approx \left[\phi''_p(x^\star)\right]^{-1}$ where $x^\star$ is the the mode of $p(x)$.

[1] For non-Gaussian approximations, the expected values of all sufficient statistics of the exponential family are equal.

2 Results

In this section, we will give tight bounds on the quality of the EP approximation (ie: of fixed-points of the EP iteration). Our results lean on the properties of log-concave distributions [10]. In section 2.1, we introduce new bounds on the moments of log-concave distributions. The bounds show that those distributions are in a certain sense close to being Gaussian. We then apply these results to study fixed points of EP, where they enable us to compute bounds on the distance between the mean and variance of the true distribution $p(x)$ and of the approximation given by EP, which we do in section 2.2.

Our bounds require us to assume that all sites $f_i(x)$ are β_m-strongly log-concave with slowly-changing log-function. That is, if we note $\phi_i(x) = -\log(f_i(x))$:

$$\forall i \; \forall x \; \phi_i^{''}(x) \geq \beta_m > 0 \tag{1}$$

$$\forall i \; \forall d \in [3,4,5,6] \; \left|\phi_i^{(d)}(x)\right| \leq K_d \tag{2}$$

The target distribution $p(x)$ then inherits those properties from the sites. Noting $\phi_p(x) = -\log(p(x)) = \sum_i \phi_i(x)$, then ϕ_p is $n\beta_m$-strongly log-concave and its higher derivatives are bounded:

$$\forall x, \; \phi_p^{''}(x) \geq n\beta_m \tag{3}$$

$$\forall d \in [3,4,5,6] \; \left|\phi_p^{(d)}(x)\right| \leq nK_d \tag{4}$$

A natural concern here is whether or not our conditions on the sites are of practical interest. Indeed, strongly-log-concave likelihoods are rare. We picked these strong regularity conditions because they make the proofs relatively tractable (although still technical and long). The proof technique carries over to more complicated, but more realistic, cases. One such interesting generalization consists of the case in which $p(x)$ and all hybrids at the fixed-point are log-concave with slowly changing log-functions (with possibly differing constants). In such a case, while the math becomes more unwieldy, similar bounds as ours can be found, greatly extending the scope of our results. The results we present here should thus be understood as a stepping stone and not as the final word on the quality of the EP approximation: we have focused on providing a rigorous but extensible proof.

2.1 Log-concave distributions are strongly constrained

Log-concave distributions have many interesting properties. They are of course unimodal, and the family is closed under both marginalization and multiplication. For our purposes however, the most important property is a result due to Brascamp and Lieb [11], which bounds their even moments. We give here an extension in the case of log-concave distributions with slowly changing log-functions (as quantified by eq. (2)). Our results show that these are close to being Gaussian.

The Brascamp-Lieb inequality states that, if $LC(x) \propto \exp(-\phi(x))$ is β_m-strongly log-concave (ie: $\phi^{''}(x) \geq \beta_m$), then centered even moments of LC are bounded by the corresponding moments of a Gaussian with variance β_m^{-1}. If we note these moments m_{2k} and $\mu_{LC} = E_{LC}(x)$ the mean of LC:

$$\begin{aligned} m_{2k} &= E_{LC}\left((x-\mu_{LC})^{2k}\right) \\ m_{2k} &\leq (2k-1)!!\beta_m^{-k} \end{aligned} \tag{5}$$

where $(2k-1)!!$ is the double factorial: the product of all odd terms from 1 to $2k-1$. $3!! = 3$, $5!! = 15$, $7!! = 105$, etc. This result can be understood as stating that a log-concave distribution must have a small variance, but doesn't generally need to be close to a Gaussian.

With our hypothesis of slowly changing log-functions, we were able to improve on this result. Our improved results include a bound on *odd* moments, as well as first order expansions of even moments (eqs. (6)-(9)).

Our extension to the Brascamp-Lieb inequality is as follows. If ϕ is slowly changing in the sense that some of its higher derivatives are bounded, as per eq. 2, then we can give a bound on $\phi^{'}(\mu_{LC})$

(showing that μ_{LC} is close to the mode $x^\star$ of LC, see eqs. (10) to (13)) and m_3 (showing that LC is mostly symmetric):

$$\left|\phi^{'}(\mu_{LC})\right| \leq \frac{K_3}{2\beta_m} \tag{6}$$

$$|m_3| \leq \frac{2K_3}{\beta_m^3} \tag{7}$$

and we can compute the first order expansions of m_2 and m_4, and bound the errors in terms of β_m and the K's :

$$\left|m_2^{-1} - \phi^{''}(\mu_{LC})\right| \leq \frac{K_3^2}{\beta_m^2} + \frac{K_4}{2\beta_m} \tag{8}$$

$$\left|\phi^{''}(\mu_{LC})m_4 - 3m_2\right| \leq \frac{19}{2}\frac{K_3^2}{\beta_m^4} + \frac{5}{2}\frac{K_4}{\beta_m^3} \tag{9}$$

With eq. (8) and (9), we see that $m_2 \approx \left(\phi^{''}(\mu_{LC})\right)^{-1}$ and $m_4 \approx 3\left(\phi^{''}(\mu_{LC})\right)^{-2}$ and, in that sense, that $LC(x)$ is close to the Gaussian with mean μ_{LC} and inverse-variance $\phi^{''}(\mu_{LC})$.

These expansions could be extended to further orders and similar formulas can be found for the other moments of $LC(x)$: for example, any odd moments can be bounded by $|m_{2k+1}| \leq C_k K_3 \beta_m^{-(k+1)}$ (with C_k some constant) and any even moment can be found to have first-order expansion: $m_{2k} \approx (2k-1)!!\left(\phi^{''}(\mu_{LC})\right)^{-k}$. The proof, as well as more detailed results, can be found in the Supplement.

Note how our result relates to the Bernstein-von Mises theorem, which says that, in the limit of a large amount of observations, a posterior $p(x)$ tends towards its CGA. If we consider the posterior obtained from n likelihood functions that are all log-concave and slowly changing, our results show the slightly different result that the moments of that posterior are *close* to those of a Gaussian with mean μ_{LC} (instead of $x^\star_{LC}$) and inverse-variance $\phi^{''}(\mu_{LC})$ (instead of $\phi^{''}(x^\star_{LC})$) . This point is critical. While the CGA still ends up capturing the limit behavior of p, as $\mu_{LC} \to x^\star$ in the large-data limit (see eq. (13) below), an approximation that would return the Gaussian approximation at μ_{LC} would be better. This is essentially what EP does, and this is how it improves on the CGA.

2.2 Computing bounds on EP approximations

In this section, we consider a given EP fixed-point $q_k^*(x|r_i, \beta_i)$ and the corresponding approximation of $p(x)$: $q^*(x|r = \sum r_i, \beta = \sum \beta_i)$. We will show that the expected value and variance of q^*(resp. μ_{EP} and v_{EP}) are close to the true mean and variance of p (resp. μ and v), and also investigate the quality of the CGA ($\mu \approx x^\star$, $v \approx \left[\phi_p^{''}(x^\star)\right]^{-1}$).

Under our assumptions on the sites (eq. (1) and (2)), we are able to derive bounds on the quality of the EP approximation. The proof is quite involved and long, and we will only present it in the Supplement. In the main text, we give a partial version: we detail the first step of the demonstration, which consists of computing a rough bound on the distance between the true mean μ, the EP approximation μ_{EP} and the mode $x^\star$, and give an outline of the rest of the proof.

Let's show that μ, μ_{EP} and $x^\star$ are all close to one another. We start from eq. (6) applied to $p(x)$:

$$\left|\phi_p^{'}(\mu)\right| \leq \frac{K_3}{2\beta_m} \tag{10}$$

which tells us that $\phi_p^{'}(\mu) \approx 0$. μ must thus be close to $x^\star$. Indeed:

$$\begin{aligned} \left|\phi_p^{'}(\mu)\right| &= \left|\phi_p^{'}(\mu) - \phi_p^{'}(x^\star)\right| \quad (11) \\ &= \left|\phi_p^{''}(\xi)(\mu - x^\star)\right| \; \xi \in [\mu, x^\star] \\ &\geq \left|\phi_p^{''}(\xi)\right| |\mu - x^\star| \\ &\geq n\beta_m |\mu - x^\star| \quad (12) \end{aligned}$$

Combining eq. (10) and (12), we finally have:

$$|\mu - x^\star| \leq n^{-1} \frac{K_3}{2\beta_m^2} \quad (13)$$

Let's now show that μ_{EP} is also close to $x^\star$. We proceed similarly, starting from eq. (6) but applied to all hybrids $h_i(x)$:

$$\forall i \; \left|\phi_i^{'}(\mu_{EP}) + \beta_{-i}\mu_{EP} - r_{-i}\right| \leq n^{-1} \frac{K_3}{2\beta_m} \quad (14)$$

which is not really equivalent to eq. (10) yet. Recall that $q(x|r,\beta)$ has mean μ_{EP}: we thus have: $r = \beta\mu_{EP}$. Which gives:

$$\begin{aligned} \left(\sum_i \beta_{-i}\right)\mu_{EP} &= ((n-1)\beta)\,\mu_{EP} \\ &= (n-1)r \\ &= \sum_i r_{-i} \quad (15) \end{aligned}$$

If we sum all terms in eq. (14), the $\beta_{-i}\mu_{EP}$ and r_{-i} thus cancel, leaving us with:

$$\left|\phi_p^{'}(\mu_{EP})\right| \leq \frac{K_3}{2\beta_m} \quad (16)$$

which is equivalent to eq. (10) but for μ_{EP} instead of μ. This shows that μ_{EP} is, like μ, close to $x^\star$:

$$|\mu_{EP} - x^\star| \leq n^{-1} \frac{K_3}{2\beta_m^2} \quad (17)$$

At this point, we can show that, since they are both close to $x^\star$ (eq. (13) and (17)), $\mu = \mu_{EP} + \mathcal{O}\left(n^{-1}\right)$, which constitutes the first step of our computation of bounds on the quality of EP.

After computing this, the next step is evaluating the quality of the approximation of the variance, via computing $\left|v^{-1} - v_{EP}^{-1}\right|$ for EP and $\left|v^{-1} - \phi_p^{''}(x^\star)\right|$ for the CGA, from eq. (8). In both cases, we find:

$$\begin{aligned} v^{-1} &= v_{EP}^{-1} + \mathcal{O}(1) \quad (18) \\ &= \phi_p^{''}(x^\star) + \mathcal{O}(1) \quad (19) \end{aligned}$$

Since v^{-1} is of order n, because of eq. (5) (Brascamp-Lieb upper bound on variance), this is a decent approximation: the relative error is of order n^{-1}.

We can find similarly that both EP and CGA do a good job of finding a good approximation of the fourth moment of p: m_4. For EP this means that the fourth moment of each hybrid and of q are a close match:

$$\begin{aligned} \forall i \; m_4 &\approx m_4^i \approx 3v_{EP}^2 \quad (20) \\ &\approx 3\left(\phi_p^{''}(m)\right)^{-2} \quad (21) \end{aligned}$$

In contrast, the third moment of the hybrids doesn't match at all the third moment of p, but their sum does !

$$m_3 \approx \sum_i m_3^i \quad (22)$$

Finally, we come back to the approximation of μ by μ_{EP}. These obey two very similar relationships:

$$\phi_p^{'}(\mu) + \phi_p^{(3)}(\mu)\frac{v}{2} = \mathcal{O}\left(n^{-1}\right) \quad (23)$$

$$\phi_p^{'}(\mu_{EP}) + \phi_p^{(3)}(\mu_{EP})\frac{v_{EP}}{2} = \mathcal{O}\left(n^{-1}\right) \quad (24)$$

Since $v = v_{EP} + \mathcal{O}\left(n^{-2}\right)$ (a slight rephrasing of eq. (18)), we finally have:

$$\mu = \mu_{EP} + \mathcal{O}\left(n^{-2}\right) \quad (25)$$

We summarize the results in the following theorem:

Theorem 1. *Characterizing fixed-points of EP*

Under the assumptions given by eq. (1) *and* (2) *(log-concave sites with slowly changing log), we can bound the quality of the EP approximation and the CGA:*

$$\begin{aligned} |\mu - x^*| &\leq n^{-1}\frac{K_3}{2\beta_m^2} \\ |\mu - \mu_{EP}| &\leq B_1(n) = \mathcal{O}\left(n^{-2}\right) \\ \left|v^{-1} - \phi_p^{''}(x^*)\right| &\leq \frac{2K_3^2}{\beta_m^2} + \frac{K_4}{2\beta_m} \\ \left|v^{-1} - v_{EP}^{-1}\right| &\leq B_2(n) = \mathcal{O}\left(1\right) \end{aligned}$$

We give the full expression for the bounds B_1 and B_2 in the Supplement

Note that the order of magnitude of the bound on $|\mu - x^\star|$ is the best possible, because it is attained for certain distributions. For example, consider a Gamma distribution with natural parameters $(n\alpha, n\beta)$ whose mean $\frac{\alpha}{\beta}$ is approximated at order n^{-1} by its mode $\frac{\alpha}{\beta} - \frac{1}{n\beta}$. More generally, from eq. (23), we can compute the first order of the error:

$$\mu - m \approx -\frac{\phi_p^{(3)}(\mu)}{\phi_p^{''}(\mu)}\frac{v}{2} \approx -\frac{1}{2}\frac{\phi_p^{(3)}(\mu)}{\left[\phi_p^{''}(\mu)\right]^2} \quad (26)$$

which is the term causing the order n^{-1} error. Whenever this term is significant, it is thus safe to conclude that EP improves on the CGA.

Also note that, since v^{-1} is of order n, the relative error for the v^{-1} approximation is of order n^{-1} for both methods. Despite having a convergence rate of the same order, the EP approximation is demonstrably better than the CGA, as we show next. Let us first see why the approximation for v^{-1} is only of order 1 for both methods. The following relationship holds:

$$v^{-1} = \phi_p^{''}(\mu) + \phi_p^{(3)}(\mu)\frac{m_3^p}{2v} + \phi_p^{(4)}(\mu)\frac{m_4^p}{3!v} + \mathcal{O}\left(n^{-1}\right) \quad (27)$$

In this relationship, $\phi_p^{''}(\mu)$ is an order n term while the rest are order 1. If we now compare this to the CGA approximation of v^{-1}, we find that it fails at multiple levels. First, it completely ignores the two order 1 terms, and then, because it takes the value of $\phi_p^{''}$ at $x^\star$ which is at a distance of $\mathcal{O}\left(n^{-1}\right)$ from μ, it adds another order 1 error term (since $\phi_p^{(3)} = \mathcal{O}\left(n\right)$). The CGA is thus adding quite a bit of error, even if each component is of order 1.

Meanwhile, v_{EP} obeys a relationship similar to eq. (27):

$$v_{EP}^{-1} = \phi_p^{''}(\mu_{EP}) + \sum_i \left[\phi_i^{(3)}(\mu_{EP})\frac{m_3^i}{2v_{EP}}\right] + \phi_p^{(4)}(\mu_{EP})\frac{3v_{EP}^2}{3!v_{EP}} + \mathcal{O}\left(n^{-1}\right) \quad (28)$$

We can see where the EP approximation produces errors. The $\phi_p^{''}$ term is well approximated: since $|\mu - \mu_{EP}| = \mathcal{O}\left(n^{-2}\right)$, we have $\phi_p^{''}(\mu) = \phi_p^{''}(\mu_{EP}) + \mathcal{O}\left(n^{-1}\right)$. The term involving m_4 is also well

approximated, and we can see that the only term that fails is the m_3 term. The order 1 error is thus entirely coming from this term, which shows that EP performance suffers more from the skewness of the target distribution than from its kurtosis.

Finally, note that, with our result, we can get some intuitions about the quality of the EP approximation using other metrics. For example, if the most interesting metric is the KL divergence $KL(p,q)$, the excess KL divergence from using the EP approximation q instead of the true minimizer q_{KL} (which has the same mean μ and variance v as p) is given by:

$$\Delta KL = \int p \log \frac{q_{KL}}{q} = \int p(x) \left(-\frac{(x-\mu)^2}{2v} + \frac{(x-\mu_{EP})^2}{2v_{EP}} - \frac{1}{2} \log \left(\frac{v}{v_{EP}} \right) \right) \quad (29)$$

$$= \frac{1}{2} \left[\frac{v}{v_{EP}} - 1 - \log \left(\frac{v}{v_{EP}} \right) \right] + \frac{(\mu-\mu_{EP})^2}{2v_{EP}} \quad (30)$$

$$\approx \frac{1}{4} \left(\frac{v-v_{EP}}{v_{EP}} \right)^2 + \frac{(\mu-\mu_{EP})^2}{2v_{EP}} \quad (31)$$

which we recognize as $KL(q_{KL}, q)$. A similar formula gives the excess KL divergence from using the CGA instead of q_{KL}. For both methods, the variance term is of order n^{-2} (though it should be smaller for EP), but the mean term is of order n^{-3} for EP while it is of order n^{-1} for the CGA. Once again, EP is found to be the better approximation.

Finally, note that our bounds are quite pessimistic: the true value might be a much better fit than we have predicted here.

A first cause is the bounding of the derivatives of $\log(p)$ (eqs. (3),(4)): while those bounds are correct, they might prove to be very pessimistic. For example, if the contributions from the sites to the higher-derivatives cancel each other out, a much lower bound than nK_d might apply. Similarly, there might be another lower bound on the curvature much higher than $n\beta_m$.

Another cause is the bounding of the variance from the curvature. While applying Brascamp-Lieb requires the distribution to have high log-curvature everywhere, a distribution with high-curvature close to the mode and low-curvature in the tails still has very low variance: in such a case, the Brascamp-Lieb bound is very pessimistic.

In order to improve on our bounds, we will thus need to use tighter bounds on the log-derivatives of the hybrids and of the target distribution, but we will also need an extension of the Brascamp-Lieb result that can deal with those cases where a distribution is strongly log-concave around its mode but, in the tails, the log-curvature is much lower.

3 Conclusion

EP has been used for now quite some time without any theoretical concrete guarantees on its performance. In this work, we provide explicit performance bounds and show that EP is superior to the CGA, in the sense of giving provably better approximations of the mean and variance. There are now theoretical arguments for substituting EP to the CGA in a number of practical problems where the gain in precision is worth the increased computational cost. This work tackled the first steps in proving that EP offers an appropriate approximation. Continuing in its tracks will most likely lead to more general and less pessimistic bounds, but it remains an open question how to quantify the quality of the approximation using other distance measures. For example, it would be highly useful for machine learning if one could show bounds on prediction error when using EP. We believe that our approach should extend to more general performance measures and plan to investigate this further in the future.

References

[1] Thomas P. Minka. Expectation Propagation for approximate Bayesian inference. In *UAI '01: Proceedings of the 17th Conference in Uncertainty in Artificial Intelligence*, pages 362–369, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1-55860-800-1. URL `http://portal.acm.org/citation.cfm?id=720257`.

[2] Malte Kuss and Carl E. Rasmussen. Assessing Approximate Inference for Binary Gaussian Process Classification. *J. Mach. Learn. Res.*, 6:1679–1704, December 2005. ISSN 1532-4435. URL http://portal.acm.org/citation.cfm?id=1194901.

[3] Hannes Nickisch and Carl E. Rasmussen. Approximations for Binary Gaussian Process Classification. *Journal of Machine Learning Research*, 9:2035–2078, October 2008. URL http://www.jmlr.org/papers/volume9/nickisch08a/nickisch08a.pdf.

[4] Guillaume Dehaene and Simon Barthelmé. Expectation propagation in the large-data limit. Technical report, March 2015. URL http://arxiv.org/abs/1503.08060.

[5] T. Minka. Divergence Measures and Message Passing. Technical report, 2005. URL http://research.microsoft.com/en-us/um/people/minka/papers/message-passing/minka-

[6] M. Seeger. Expectation Propagation for Exponential Families. Technical report, 2005. URL http://people.mmci.uni-saarland.de/~{}mseeger/papers/epexpfam.pdf.

[7] Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer, 1st ed. 2006. corr. 2nd printing 2011 edition, October 2007. ISBN 0387310738. URL http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/038731073

[8] Jack Raymond, Andre Manoel, and Manfred Opper. Expectation propagation, September 2014. URL http://arxiv.org/abs/1409.6179.

[9] Anirban DasGupta. *Asymptotic Theory of Statistics and Probability (Springer Texts in Statistics)*. Springer, 1 edition, March 2008. ISBN 0387759700. URL http://www.amazon.com/exec/obidos/redirect?tag=citeulike07-20&path=ASIN/03877597

[10] Adrien Saumard and Jon A. Wellner. Log-concavity and strong log-concavity: A review. *Statist. Surv.*, 8:45–114, 2014. doi: 10.1214/14-SS107. URL http://dx.doi.org/10.1214/14-SS107.

[11] Herm J. Brascamp and Elliott H. Lieb. Best constants in young's inequality, its converse, and its generalization to more than three functions. *Advances in Mathematics*, 20 (2):151–173, May 1976. ISSN 00018708. doi: 10.1016/0001-8708(76)90184-5. URL http://dx.doi.org/10.1016/0001-8708(76)90184-5.

Supplementary information of "Bounding errors of Expectation-Propagation"

A Improving on the Brascamp-Lieb bound

In this section, we detail our mathematical results concerning the extension of the Brascamp-Lieb bound.

We will note $LC(x) = \exp(-\phi(x))$ a log-concave distribution. We assume that ϕ is strongly convex, and slowly changing, ie:

$$\forall x \; \phi^{''}(x) \geq \beta_m \tag{32}$$

$$\forall d \in [3,4,5,6] \; \left|\phi^{(d)}(x)\right| \leq K_d \tag{33}$$

A.1 The original Brascamp-Lieb theorem

Let $\mu_{LC} = E_{LC}(x)$ be the expected value of LC. The original Brascamp-Lieb result [1976] concerns bounding fractional centered moments of LC by the corresponding fractional moments of a Gaussian of variance β_m^{-1}, centered at μ_{LC}. Noting $g(x) = \mathcal{N}\left(x|\mu_{LC}, \beta_m^{-1}\right)$ that Gaussian, we have:

$$\forall \alpha \geq 1 \; E_{LC}\left(|x-\mu_{LC}|^{\alpha}\right) \leq E_g\left(|x-\mu_{LC}|^{\alpha}\right) \tag{34}$$

However, we are not interested in their full result, but only in a restricted version of it which only concerns even moments. This version simply reads:

$$\forall k \in \mathbb{N} \; m_{2k} = E_{LC}\left(|x-\mu_{LC}|^{2k}\right) \leq (2k-1)\, m_{2k-2}\beta_m^{-1} \tag{35}$$

$$m_{2k} \leq (2k-1)!!\beta_m^{-k} \tag{36}$$

where $(2k-1)!!$ is the double-factorial: the product of all odd terms between 1 and $2k-1$. Eq. (35) might be a new result. Note that equality only occurs when $f(x) = 1$ and LC is Gaussian. Note also that the bounds on the higher derivatives of ϕ are not needed for this result, but only for our extension.

We offer here a proof of eq. (35) (from which eq. (36) is a trivial consequence), which is slightly different from Brascamp & Lieb's original proof. We believe this proof to be original, though it is still quite similar to the original proof.

Proof. Let's decompose $LC(x)$ into two parts:

- $g(x) = \mathcal{N}\left(x|\mu_{LC}, \beta_m^{-1}\right)$ the bounding Gaussian with same mean as LC
- $f(x) = \frac{LC(x)}{g(x)}$ the remainder

f is easily shown to be log-concave, which means that it is unimodal. We will note $x^\star$ the mode of f. f is increasing on $]-\infty, x^\star]$ and decreasing on $[x^\star, \infty[$. We thus know the sign of $f^{'}(x)$:

$$\text{sign}\left(f^{'}(x)\right) = \text{sign}\left(x^\star - x\right) \tag{37}$$

Consider the integral: $\int_{-\infty}^{+\infty} g(x)f^{'}(x)dx$. By integration by parts (or by Stein's lemma), we have:

$$\begin{aligned} \int_{-\infty}^{+\infty} g(x)f^{'}(x)dx &= \int_{-\infty}^{+\infty} g(x)f(x)\beta_m(x-\mu)dx \\ &= \beta_m(\mu-\mu) \\ \int_{-\infty}^{+\infty} g(x)f^{'}(x)dx &= 0 \end{aligned} \tag{38}$$

We now split the integral at μ_{LC} and $x^\star$, assuming without loss of generality that $x^\star \leq \mu_{LC}$:

$$\begin{aligned}\int_{-\infty}^{x^\star} gf^{'} + \int_{x^\star}^{\mu_{LC}} gf^{'} + \int_{\mu_{LC}}^{\infty} gf^{'} &= 0 \\ \int_{-\infty}^{x^\star} gf^{'} + \int_{x^\star}^{\mu_{LC}} gf^{'} &= -\int_{\mu_{LC}}^{\infty} gf^{'} \\ &\geq 0 \end{aligned} \tag{39}$$

Now consider a statistic $S_k(x) = (x - \mu_{LC})^{2k-1}$. Again using integration by parts, we have the following equality:

$$\begin{aligned}\int g(x)f^{'}(x)S_k(x)dx &= \int g(x)f(x)\left(\beta_m S_k(x)(x-\mu) - S_k^{'}(x)\right)dx \\ \int g(x)f^{'}(x)\left(x-\mu_{LC}\right)^{2k-1}dx &= \int LC(x)\left(\beta_m\left(x-\mu_{LC}\right)^{2k} - (2k-1)\left(x-\mu_{LC}\right)^{2k-2}\right) \\ &= \beta_m m_{2k} - (2k-1)m_{2k-2} \end{aligned} \tag{40}$$

At this point, we only need to prove that $\int gf^{'}S_k \leq 0$ to finish our proof, from eq. (40). We will actually prove a slightly stronger result: that even if we cut the integral at μ_{LC}, both halves are still negative:

$$\int_{-\infty}^{\mu_{LC}} gf^{'}S_k \leq 0 \tag{41}$$

$$\int_{\mu_{LC}}^{\infty} gf^{'}S_k \leq 0 \tag{42}$$

Eq. (42) is trivial. g is positive everywhere, while $S_k(x) \geq 0$ and $f^{'}(x) \leq 0$ for $x \geq \mu_{LC}$.

Eq. (41) is slightly harder. From eq. (39), $\int_{-\infty}^{x^\star} gf^{'} + \int_{x^\star}^{\mu_{LC}} gf^{'} \geq 0$, where the first term is positive, and the second negative. When we multiply the integrand by the decreasing positive function $-S_k(x) = -\left(x-\mu_{LC}\right)^{2k-1}$, the order in the terms is preserved. To say it in equations:

$$\begin{aligned}\int_{-\infty}^{x^\star} gf^{'}\left(-S_k\right) &\geq \left(-\left(x^\star-\mu_{LC}\right)^{2k-1}\right)\int_{-\infty}^{x^\star} gf^{'} \\ &\geq \left(-\left(x^\star-\mu_{LC}\right)^{2k-1}\right)\left(-\int_{x^\star}^{\mu_{LC}} gf^{'}\right) \\ &\geq \int_{x^\star}^{\mu_{LC}} gf^{'}S_k \end{aligned} \tag{43}$$

from which we finally find eq. (41), which concludes our proof. Note that there is the equality $\int gf^{'}S_k = 0$ IFF $f^{'}(x) = 0$, justifying our earlier comment about $m_{2k} = (2k-1)\beta_m^{-1}m_{2k-2}$ IFF $LC(x) = g(x)$. □

A.2 Extending the Brascamp-Lieb theorem

The original Brascamp-Lieb result tells us that the spread of $LC(x)$ (as measured by its even moments) can't be too important, but it doesn't tell us whether such distributions are close to being Gaussian, which is what EP requires. By constraining the higher derivatives of $\phi(x)$, we are able to constrain how far LC is from a Gaussian distribution. This is the essence of our extension of the Brascamp-Lieb theorem. We derived the following:

Theorem 2. *Extension of the Brascamp-Lieb theorem*

With LC a strongly log-concave distribution with slowly changing log-function (eqs. (32), (33)*), we have the following inequalities:*

$$\left|\phi^{'}\left(\mu_{LC}\right)\right| \leq \frac{K_3}{2\beta_m} \tag{44}$$

$$\left|\frac{m_3}{m_2}\right| \leq 2\frac{K_3}{\beta_m^2} \tag{45}$$

$$\left|\frac{m_5}{m_2}\right| \leq \frac{17K_3}{\beta_m^3} \tag{46}$$

which generalizes to:

$$\left|\frac{m_{2k+1}}{m_2}\right| \leq C_k \frac{K_3}{\beta_m^{k+1}} \tag{47}$$

The following first order expansions of m_2, m_3 and m_4:

$$\left|m_2^{-1} - \phi_2^{''}(\mu_{LC})\right| \leq \frac{K_3^2}{\beta_m^2} + \frac{K_4}{2\beta_m} \tag{48}$$

$$\left|\phi^{''}(\mu_{LC})m_2 - 1\right| \leq \frac{K_3^2}{\beta_m^3} + \frac{K_4}{2\beta_m^2} \tag{49}$$

$$\left|\phi^{''}(\mu_{LC})m_3 + \left(\phi^{'}(\mu_{LC})m_2 + \frac{\phi^{(3)}(\mu_{LC})}{2}m_4\right)\right| \leq \frac{17}{6}\frac{K_3K_4}{\beta_m^4} + \frac{5}{8}\frac{K_5}{\beta_m^3} \tag{50}$$

$$\left|\phi^{''}(\mu_{LC})m_4 - 3m_2\right| \leq \frac{19}{2}\frac{K_3^2}{\beta_m^4} + \frac{5}{2}\frac{K_4}{\beta_m^3} \tag{51}$$

which generalizes to:

$$m_{2k+2} \approx \frac{(2k+1)}{\phi^{''}(\mu_{LC})}m_{2k} \tag{52}$$

$$\approx (2k+1)!!\left[\phi^{''}(\mu_{LC})\right]^{-(k+1)} \tag{53}$$

And the following higher order relationships:

$$\left|\phi^{'}(\mu_{LC}) + \frac{\phi^{(3)}(\mu_{LC})}{2}m_2\right| \leq \frac{K_3K_4}{3\beta_m^3} + \frac{K_5}{8\beta_m^2} \tag{54}$$

$$\left|m_2^{-1} - \phi^{''}(\mu_{LC}) - \frac{\phi^{(3)}(\mu_{LC})}{2}\frac{m_3}{m_2} - \frac{\phi^{(4)}(\mu_{LC})}{3!}\frac{m_4}{m_2}\right| \leq \frac{17}{24}\frac{K_3K_5}{\beta_m^3} + \frac{K_6}{8\beta_m^2} \tag{55}$$

Note that we refer to eq. (48), (50) and (51) as first order expansions because you can read them as, respectively:

$$\begin{aligned}
m_2 &\approx \left(\phi^{''}(\mu_{LC})\right)^{-1} \\
m_3 &\approx -\left(\phi^{''}(\mu_{LC})\right)^{-1}\left(\phi^{'}(\mu_{LC})m_2 + \frac{\phi^{(3)}(\mu_{LC})}{2}m_4\right) \\
m_4 &\approx 3\left(\phi^{''}(\mu_{LC})\right)^{-1}m_2
\end{aligned}$$

These relationships are not exhaustive, and one could find many such relationships for even higher orders. The list presented here only concerns results which we will need for our bound on EP.

Proof. We will first give an outline of the proof, and then dive into all the equations of the full proof.

The key component of the proof is Stein's lemma (ie: integration by parts). For $LC = \exp(-\phi(x))$, it reads: for any statistic $S(x)$ with at-most-polynomial growth:

$$E_{LC}\left(\phi^{'}(x)S(x) - S^{'}(x)\right) = 0 \tag{56}$$

which we will only use for statistics of the form $S_k(x) = (x - \mu_{LC})^k$. This gives us the following relationships:

$$\begin{aligned} E_{LC}\left(\phi^{'}(x)\right) &= 0 & (57)\\ E_{LC}\left(\phi^{'}(x)(x-\mu_{LC})\right) &= 1 & (58)\\ E_{LC}\left(\phi^{'}(x)(x-\mu_{LC})^2\right) &= 0 & (59)\\ E_{LC}\left(\phi^{'}(x)(x-\mu_{LC})^3\right) &= 3m_2 & (60) \end{aligned}$$

and further relationships of the same form that we won't need. The key intuition in understanding why LC is almost Gaussian is the following: $\phi^{'}(x) \approx \phi^{''}(\mu_{LC})(x-\mu)$. The Stein relationships for LC are thus almost the same relationships that would be obeyed by the Gaussian $g_{\mu_{LC}}(x) = \mathcal{N}\left(x|\mu_{LC}, \left(\phi^{''}(\mu_{LC})\right)^{-1}\right)$. This is why LC is close to $g_{\mu_{LC}}$.

For all these relationships, we will perform a Taylor expansion around μ_{LC}, which now gives us self-consistency relationships between the different moments of LC. For example, just keeping the first term in eq. (57) gives us eq. (44):

$$\phi^{'}(\mu_{LC}) \approx 0$$

We need to be careful with how we deal with the remainder of the Taylor approximation. Using the Taylor-Lagrange formula, we can bound the error that results from cutting off the Taylor series after some term, with a term of the form $C \times (x-\mu)^k$ for some constant C. The expected value under LC of that term can then bounded from the Brascamp-Lieb theorem. For example, to perform the cut-off of eq. (57) we just did, we start from the Taylor-Lagrange expression:

$$\left|\phi^{'}(x) - \phi^{'}(\mu_{LC}) - \phi^{''}(\mu_{LC})(x-\mu_{LC})\right| \leq \frac{K_3}{2}(x-\mu_{LC})^2 \tag{61}$$

which, when we take the expected value, becomes:

$$\left|E_{LC}\left(\phi^{'}(x)\right) - \phi^{'}(\mu_{LC})\right| \leq \frac{K_3}{2}m_2 \leq \frac{K_3}{2\beta_m} \tag{62}$$

where we have applied the Brascamp-Lieb theorem. This concludes the proof of eq. (44), and our introduction to the full proof.

Let's now prove the second relationship of the theorem: eq. (45). We start from eq. (59). We perform the expansion of $\phi^{'}(x)$ up to the $\phi^{''}(\mu_{LC})(x-\mu_{LC})$ term. From Taylor-Lagrange, the error is:

$$\begin{aligned} \left|\phi^{'}(x) - \phi'(\mu_{LC}) - \phi^{''}(\mu_{LC})(x-\mu_{LC})\right| &\leq \frac{K_3}{2}(x-\mu_{LC})^2\\ \left|\phi^{'}(x)(x-\mu_{LC})^2 - \phi'(\mu_{LC})(x-\mu_{LC})^2 - \phi^{''}(\mu_{LC})(x-\mu_{LC})^3\right| &\leq \frac{K_3}{2}(x-\mu_{LC})^4 \end{aligned} \tag{63}$$

We now take the expected value:

$$\left|\phi'(\mu_{LC})m_2 + \phi^{''}(\mu_{LC})m_3 - E_{LC}\left(\phi^{'}(x)\right)\right| \leq \frac{K_3}{2}m_4 \tag{64}$$

Finally, we divide by m_2, take out the $\phi^{'}(\mu_{LC})$ term from the absolute value, use the bound on $\frac{m_4}{m_2}$ from eq. (35), and lower bound $\phi^{''}(\mu_{LC})$:

$$\begin{aligned}
\left|\phi^{''}(\mu_{LC})\frac{m_3}{m_2}\right| &\leq \frac{K_3}{2}\frac{m_4}{m_2} + \left|\phi^{'}(\mu_{LC})\right| \\
&\leq \frac{K_3}{2\beta_m}(3) + \frac{K_3}{2\beta_m} \\
&\leq \frac{2K_3}{\beta_m} \qquad (65) \\
\left|\frac{m_3}{m_2}\right| &\leq \frac{2K_3}{\beta_m^2} \qquad (66)
\end{aligned}$$

which gives us eq. (45).

Now, let's prove the bound on m_5 (eq. (46)). The demonstration is quite similar to the m_3 bound. We start from another Stein relationship:

$$E_{LC}\left(\phi^{'}(x)\left(x-\mu_{LC}\right)^4\right) = 4m_3$$

With the same Taylor-Lagrange expansion as in eq. (63) and after taking the expected value, we have:

$$\left|4m_3 - \phi^{'}(\mu_{LC})m_4 - \phi^{''}(\mu_{LC})m_5\right| \leq \frac{K_3}{2}m_6 \qquad (67)$$

Which we divide by m_2 and manipulate further:

$$\begin{aligned}
\left|\phi^{''}(\mu_{LC})\frac{m_5}{m_2}\right| &\leq 4\left|\frac{m_3}{m_2}\right| + \left|\phi^{'}(\mu_{LC})\right|\frac{m_4}{m_2} + \frac{K_3}{2}\frac{m_6}{m_2} \\
&\leq \frac{8K_3}{\beta_m^2} + \frac{K_3}{2\beta_m}\frac{3}{\beta_m} + \frac{K_3}{2}\frac{15}{\beta_m^2} \\
&\leq \frac{17K_3}{\beta_m^2} \\
\left|\frac{m_5}{m_2}\right| &\leq \frac{17K_3}{\beta_m^3} \qquad (68)
\end{aligned}$$

which gives us eq. (46).

In order to show that any odd centered moment admits a similar bound (as we mention it the main text), we proceed by induction. The Stein relationships:

$$E\left(\phi^{'}(x)\left(x-\mu_{LC}\right)^{2k} - 2k\left(x-\mu_{LC}\right)^{2k-1}\right) = 0$$

give us the inductive step through steps identical to the preceeding equations, and we have already have the initialization (from eq. 66). We can thus find similar bounds for any higher odd moment of $LC(x)$.

Now we will prove the first order expansions, starting with the one for m_2 (eq. (48)). We now start from eq. (58), which is:

$$E_{LC}\left(\phi^{'}(x)\left(x-\mu_{LC}\right)\right) = 1$$

First step, the Taylor-Lagrange expansion. We cut off the Taylor series at $\frac{\phi^{(3)}(\mu_{LC})}{2}\left(x-\mu_{LC}\right)^2$. We can bound the error with:

$$\left|\phi^{'}(x)\left(x-\mu_{LC}\right) - \phi^{'}(\mu_{LC})\left(x-\mu_{LC}\right) - \phi^{''}(\mu_{LC})\left(x-\mu_{LC}\right)^2 - \frac{\phi^{(3)}(\mu_{LC})}{2}\left(x-\mu_{LC}\right)^3\right| \leq \frac{K_4}{3!}\left(x-\mu_{LC}\right)^4 \qquad (69)$$

which becomes, when we take the expected value:

$$\begin{aligned}
\left|1 - 0 - \phi^{''}(\mu_{LC})m_2 - \frac{\phi^{(3)}(\mu_{LC})}{2}m_3\right| &\leq \frac{K_4}{3!}m_4 \\
\left|\phi^{''}(\mu_{LC})m_2 - 1\right| &\leq \frac{1}{2}\left|\phi^{(3)}(\mu_{LC})\right||m_3| + \frac{K_4}{3!}m_4 \\
\left|m_2^{-1} - \phi^{''}(\mu_{LC})\right| &\leq \frac{1}{2}\left|\phi^{(3)}(\mu_{LC})\right|\left|\frac{m_3}{m_2}\right| + \frac{K_4}{3!}\frac{m_4}{m_2} \qquad (70) \\
&\leq \frac{K_3^2}{\beta_m^2} + \frac{K_4}{2\beta_m} \qquad (71)
\end{aligned}$$

which proves eq. (48), from which eq. (49) is a trivial consequence.

Now, the m_3 first order expansion (eq. (50)). We start from the Stein relationship from eq. (59) (which we already used to prove the bound on $\left|\frac{m_3}{m_2}\right|$).

$$E_{LC}\left(\phi^{'}(x)\left(x - \mu_{LC}\right)^2\right) = 0$$

The difference between the m_3 bound and the m_3 first order expansion is that we take a higher-order expansion of $\phi^{'}(x)$. This time, we stop at $\phi^{(4)}(\mu_{LC})\left(x - \mu_{LC}\right)^3$. The Taylor-Lagrange error is bounded by $\frac{K_5}{4!}\left(x - \mu_{LC}\right)^4$. This gives us the following bound once we take the expected value.

$$\left|\phi^{'}(\mu_{LC})m_2 + \phi^{''}(\mu_{LC})m_3 + \frac{\phi^{(3)}(\mu_{LC})}{2}m_4 + \frac{\phi^{(4)}(\mu_{LC})}{3!}m_5\right| \leq \frac{K_5}{4!}m_6 \qquad (72)$$

In that equation, m_5 is an order of magnitude smaller than the other terms, and we take it out of the absolute value:

$$\begin{aligned}
\left|\phi^{''}(\mu_{LC})m_3 - \left(-\phi'(\mu_{LC})m_2 - \frac{\phi^{(3)}(\mu_{LC})}{2}m_4\right)\right| &\leq \frac{\left|\phi^{(4)}(\mu_{LC})\right|}{3!}|m_5| + \frac{K_5}{4!}m_6 \\
&\leq \frac{K_4}{3!}\frac{17K_3}{\beta_m^4} + \frac{K_5}{4!}\frac{15}{\beta_m^3} \\
&\leq \frac{17}{6}\frac{K_3K_4}{\beta_m^4} + \frac{5}{8}\frac{K_5}{\beta_m^3} \qquad (73)
\end{aligned}$$

which proves eq. (50).

Finally, we prove the last first order expansion: eq. (51) concerning m_4. We start from the last Stein relationship: eq. (60):

$$E_{LC}\left(\phi^{'}(x)\left(x - \mu_{LC}\right)^3\right) = 3m_2$$

We cut-off the Taylor series after $\frac{\phi^{(3)}(\mu_{LC})}{2}\left(x - \mu_{LC}\right)^2$. After taking the expected value, the error is:

$$\left|3m_2 - \phi^{'}(\mu_{LC})m_3 - \phi^{''}(\mu_{LC})m_4 - \frac{\phi^{(3)}(\mu_{LC})}{2}m_5\right| \leq \frac{K_4}{3!}m_6 \qquad (74)$$

In this expression, $\phi^{'}(\mu_{LC})m_3$ and $\frac{\phi^{(3)}(\mu_{LC})}{2}m_5$ are both smaller by an order of magnitude, and we remove them from the absolute value, to finally obtain:

$$\begin{aligned}
\left|\phi^{''}(\mu_{LC})m_4 - 3m_2\right| &\leq \left|\phi^{'}(\mu_{LC})\right||m_3| + \left|\frac{\phi^{(3)}(\mu_{LC})}{2}\right||m_5| + \frac{K_4}{3!}m_6 \\
&\leq \frac{K_3}{2\beta_m}\frac{2K_3}{\beta_m^3} + \frac{K_3}{2}\frac{17K_3}{\beta_m^4} + \frac{K_4}{3!}\frac{15}{\beta_m^3} \\
&\leq \frac{19}{2}\frac{K_3^2}{\beta_m^4} + \frac{5}{2}\frac{K_4}{\beta_m^3} \qquad (75)
\end{aligned}$$

which proves eq. (51).

In order to find the first order developments of higher order even moments, one proceeds identically to here but from the Stein relationships:

$$E\left(\phi^{'}(x)\left(x-\mu_{LC}\right)^{2k+1}-(2k+1)\left(x-\mu_{LC}\right)^{2k}\right) \tag{76}$$

from which, by the same approach as the proof of eq. 51, we have:

$$m_{2k+2} \approx \frac{(2k+1)}{\phi^{''}(\mu_{LC})} m_{2k} \tag{77}$$

and by induction, we prove that:

$$m_{2k+2} \approx (2k+1)!! \left[\phi^{''}(\mu_{LC})\right]^{-(k+1)} \tag{78}$$

which justifies our claim in the main text.

We are only left with proving the final two relationships. For eq. (54), this corresponds to doing a further expansion of the first Stein relationship (eq. (57), from which we proved that $\phi^{'}(\mu_{LC}) \approx 0$):

$$E_{LC}\left(\phi^{'}(x)\right) = 0$$

We stop the Taylor series after $\frac{\phi^{(4)}(\mu_{LC})}{3!}\left(x-\mu_{LC}\right)^3$. After taking the expected value, we get:

$$\left|\phi^{'}(\mu_{LC}) + \frac{\phi^{(3)}(\mu_{LC})}{2} m_2 + \frac{\phi^{(4)}(\mu_{LC})}{3!} m_3\right| \leq \frac{K_5}{4!} m_4 \tag{79}$$

We extract the m_3 term which is an order of magnitude smaller than the other ones, and obtain:

$$\begin{aligned}
\left|\phi^{'}(\mu_{LC}) + \frac{\phi^{(3)}(\mu_{LC})}{2} m_2\right| &\leq \left|\frac{\phi^{(4)}(\mu_{LC})}{3!}\right| |m_3| + \frac{K_5}{4!} m_4 \\
&\leq \frac{K_4}{3!}\frac{2K_3}{\beta_m^3} + \frac{K_5}{4!}\frac{3}{\beta_m^2} \\
&\leq \frac{K_3 K_4}{3\beta_m^3} + \frac{K_5}{8\beta_m^2}
\end{aligned} \tag{80}$$

which proves eq. (54).

At last, we reach the proof of eq. (55). We start from the second Stein relationship (eq. (58), which we already used to get the first order expansion of m_2):

$$E_{LC}\left(\phi^{'}(x)\left(x-\mu_{LC}\right)\right) = 1$$

We stop the Taylor series after $\frac{\phi^{(5)}(\mu_{LC})}{4!}\left(x-\mu_{LC}\right)^4$. After taking the expected value, we get:

$$\left|1 - \phi^{''}(\mu_{LC}) m_2 - \frac{\phi^{(3)}(\mu_{LC})}{2} m_3 - \frac{\phi^{(4)}(\mu_{LC})}{3!} m_4 - \frac{\phi^{(5)}(\mu_{LC})}{4!} m_5\right| \leq \frac{K_6}{5!} m_6 \tag{81}$$

We divide by m_2, then extract the m_5 term and obtain:

$$\begin{aligned}
\left|m_2^{-1} - \phi^{''}(\mu_{LC}) + \frac{\phi^{(3)}(\mu_{LC})}{2}\frac{m_3}{m_2} + \frac{\phi^{(4)}(\mu_{LC})}{3!}\frac{m_4}{m_2}\right| &\leq \left|\frac{\phi^{(5)}(\mu_{LC})}{4!}\right|\left|\frac{m_5}{m_2}\right| + \frac{K_6}{5!}\frac{m_6}{m_2} \\
&\leq \frac{K_5}{4!}\frac{17K_3}{\beta_m^3} + \frac{K_6}{5!}\frac{15}{\beta_m^2} \\
&\leq \frac{17}{24}\frac{K_3 K_5}{\beta_m^3} + \frac{K_6}{8\beta_m^2}
\end{aligned} \tag{82}$$

proving eq. (55) and concluding our proof. □

B Quality of fixed-points of EP

In this section, we give a detailed proof of our bounds on the quality of the EP approximation.

We assume that all sites $f_i = \exp(-\phi_i(x))$ are β_m-strongly log-concave, with slowly changing log-functions. That is:

$$\begin{aligned} \forall i, x,\ \phi_i^{''}(x) &\geq \beta_m & (83) \\ \forall d \in [3,4,5,6]\ \left|\phi_d^{(3)}(x)\right| &\leq K_d & (84) \end{aligned}$$

The target distribution $p(x)$ then inherits those properties from the sites. Noting $\phi_p(x) = -\log(p(x)) = \sum_i \phi_i(x)$, then ϕ_p is $n\beta_m$-strongly log-concave and for $d \in [3,4,5,6]$,

$$\left|\phi_p^{(d)}(x)\right| \leq nK_d \tag{85}$$

Let $q_i(x|r_i,\beta_i)$ be the site-approximations of a fixed-point of EP, $q(x|r = \sum_i r_i, \beta = \sum_i \beta_i)$ be the corresponding approximation of $p(x)$ and $h_i(x)$ the corresponding hybrid distributions. From our hypothesis on the sites, all hybrids are $(\beta_m + \beta_{-i})$-strongly log-concave, with slowly varying log-function (with constants K_d). We can thus apply our results from section A to all hybrids and the target distribution.

Some results to keep in mind on the hybrids: first of all,

$$-\frac{\partial \log(h_i(x))}{\partial x} = \phi_i^{'}(x) + \beta_{-i}x - r_{-i} \tag{86}$$

This expression is important as it is the one that appears in the Stein relationships.

Also, because $q(x)$ is a Gaussian distribution of mean and variance μ_{EP}, v_{EP} and with natural parameters r, β:

$$\begin{aligned} r &= \beta\mu_{EP} & (87) \\ \beta &= v_{EP}^{-1} & (88) \end{aligned}$$

Finally, we have:

$$\begin{aligned} \sum_i \beta_{-i}\mu_{EP} &= \sum_{i,j\neq i} \beta_j \mu_{EP} \\ &= (n-1)\sum_j \beta_j \mu_{EP} \\ &= (n-1)\,\beta\mu_{EP} \\ &= (n-1)\,r \\ &= (n-1)\sum_j r_j \\ &= \sum_{i,j\neq i} r_j \\ \sum \beta_{-i}\mu_{EP} &= \sum_i r_{-i} \end{aligned} \tag{89}$$

B.1 Lower-bounding the β_i

Let's show that we can lower bound the β_i at the fixed-point by β_m.

Recall that β_i is obtained from the difference between the inverse variance of $h_i(x)$ and β_{-i}, and $h_i(x)$ happens to be a $(\beta_m + \beta_{-i})$-strongly log-concave distribution. We can thus apply the

Brascamp-Lieb inequality to the variance:

$$m_2^i \leq \frac{1}{\beta_m + \beta_{-i}} \tag{90}$$

$$\left(m_2^i\right)^{-1} \geq \beta_m + \beta_{-i} \tag{91}$$

Thus, $\beta_i = \left(m_2^i\right)^{-1} - \beta_{-i} \geq \beta_m$ and we have the claimed lower bound.[2]

Thus all hybrids are actually at least $n\beta_m$-strongly log-concave (but could theoretically be stronger. This is one way our bounds can be pessimistic).

B.2 Approximation of various moments by $q(x)$ and the hybrids

In this section, we will show that some moments of $p(x)$ are matched approximately by the moments of $q(x)$ and/or the moments of the hybrids $h_i(x)$.

We will note m_k^p the k^{th} centered moment of $p(x)$ and m_k^i the moments of the hybrids. We will use μ, v for the mean and variance of $p(x)$ and μ_{EP}, v_{EP} for the mean and variance of $q(x)$ and all $h_i(x)$ (recall that, at a fixed-point of EP, $q(x)$ and all $h_i(x)$ share the same mean and variance). The mean and variance have gained special notation due to their special status.

With these notations, the first three even moments of q are respectively v_{EP}, $3v_{EP}^2$ and $15v_{EP}^3$, while all odd moments are 0.

We will show that the following moments are matched:

Theorem 3. *When all sites are strongly log-concave with slowly changing log, fixed-points of EP provide a good approximation of several moments of $p(x)$:*

$$\begin{aligned}
\mu &= \mu_{EP} + \mathcal{O}\left(n^{-2}\right) \\
v^{-1} &= v_{EP}^{-1} + \mathcal{O}\left(1\right) \\
m_3^p &= \sum_i m_3^i + \mathcal{O}\left(n^{-3}\right) \\
m_4^p &= 3v_{EP}^2 + \mathcal{O}\left(n^{-3}\right) \\
\forall i \; m_4^p &= m_4^i + \mathcal{O}\left(n^{-3}\right)
\end{aligned}$$

Proof. Let's first give an outline of the proof.

The logic for all these results is similar. Because all hybrids $h_i(x)$ are $n\beta_m$-strongly log-concave with slowly changing-log, we can apply the results of section A on all those distributions, and obtain inequalities that relate the moments of the $h_i(x)$ to one another. Since they all share the same mean and variance, these become severely constrained. Since $p(x)$ is also log-concave with slowly changing log-function, its mean and variance obey very similar relationships to μ_{EP} and v_{EP}. From the fact that the pair (μ, v) and the pair (μ_{EP}, v_{EP}) obey almost the same inequalities, we are able to deduce that they are close to one another.

Let's start with μ. From eq. (54), μ obeys the following simple relationship:

$$\begin{aligned}
\left|\phi_p^{'}(\mu) + \frac{\phi_p^{(3)}(\mu)}{2} v\right| &\leq \frac{nK_3 nK_4}{3n^3\beta_m^3} + \frac{nK_5}{8n^2\beta_m^2} \\
&\leq n^{-1}\left(\frac{K_3 K_4}{3\beta_m^3} + \frac{K_5}{8\beta_m^2}\right)
\end{aligned} \tag{92}$$

Applying the same results to all hybrids $h_i(x)$, we get:

[2] By the the same logic, if all sites are strongly log-concave, the dynamics of EP must always maintain $\beta_i \geq \beta_m$. It is thus useless to initialize the EP algorithm at a lower value.

$$\forall i \quad \left|\phi_i^{'}(\mu_{EP}) + \beta_{-i}\mu_{EP} - r_{-i} + \frac{\phi_i^{(3)}(\mu_{EP})}{2}v_{EP}\right| \leq \frac{K_3K_4}{3n^3\beta_m^3} + \frac{K_5}{8n^2\beta_m^2}$$
$$\leq n^{-3}\frac{K_3K_4}{3\beta_m^3} + n^{-2}\frac{K_5}{8\beta_m^2} \quad (93)$$

which is slightly different than eq. (92). Let's now sum the relationship obtained for each $h_i(x)$. The $\beta_{-i}\mu_{EP} - r_{-i}$ terms drop out (eq. (89)) and we get:

$$\left|\phi_p^{'}(\mu_{EP}) + \frac{\phi_p^{(3)}(\mu_{EP})}{2}v_{EP}\right| \leq n^{-2}\frac{K_3K_4}{3\beta_m^3} + n^{-1}\frac{K_5}{8\beta_m^2} \quad (94)$$

We have that μ and μ_{EP} satisfy almost the same relationship from eq. (92) and (94). We can use this to bound the distance between the two, as a function of the distance between v and v_{EP}:

$$\phi_p^{'}(\mu_{EP}) + \frac{\phi_p^{(3)}(\mu_{EP})}{2}v_{EP} = \phi_p^{'}(\mu_{EP}) + \frac{\phi_p^{(3)}(\mu_{EP})}{2}(v + v_{EP} - v)$$
$$\phi_p^{'}(\mu_{EP}) + \frac{\phi_p^{(3)}(\mu_{EP})}{2}v_{EP} - \left(\phi_p^{'}(\mu) + \frac{\phi_p^{(3)}(\mu)}{2}v\right) = \left[\phi_p^{''}(\xi_1) + \frac{\phi_p^{(4)}(\xi_2)}{2}v\right](\mu_{EP} - \mu)$$
$$+\frac{\phi_p^{(3)}(\mu_{EP})}{2}(v_{EP} - v) \quad (95)$$

where $\xi_1, \xi_2 \in [\mu, \mu_{EP}]$ and we have used first-order expansions at μ of $\phi_p^{'}(\mu_{EP})$ and $\phi_p^{(3)}(\mu_{EP})$. We can go from upper bounding $\left[\phi_p^{''}(\xi_1) + \frac{\phi_p^{(4)}(\xi_2)}{2}v\right](\mu_{EP} - \mu)$ to upper bounding $|\mu - \mu_{EP}|$:

$$\left|\left[\phi_p^{''}(\xi_1) + \frac{\phi_p^{(4)}(\xi_2)}{2}v\right](\mu_{EP} - \mu)\right| \geq \min_{\xi_1,\xi_2}\left(\left[\phi_p^{''}(\xi_1) + \frac{\phi_p^{(4)}(\xi_2)}{2}v\right]\right)|\mu - \mu_{EP}|$$
$$\geq \left[n\beta_m - \frac{K_4}{2\beta_m}\right]|\mu - \mu_{EP}| \quad (96)$$

We finally obtain a bound on the distance between μ and μ_{EP} by combining eqs. (92), (94), (95) and (96):

$$\left|\left[\phi_p^{''}(\xi_1) + \frac{\phi_p^{(4)}(\xi_2)}{2}v\right](\mu_{EP} - \mu)\right| \leq \left(n^{-1} + n^{-2}\right)\frac{K_3K_4}{3\beta_m^3} + 2n^{-1}\frac{K_5}{8\beta_m^2} + n\frac{K_3}{2}|v - v_{EP}|$$
$$\leq \mathcal{O}\left(n^{-1}\right) + \mathcal{O}\left(n^{-1}\right) + \mathcal{O}\left(n|v - v_{EP}|\right) \quad (97)$$
$$|\mu - \mu_{EP}| \leq \mathcal{O}\left(n^{-2}\right) + \mathcal{O}\left(|v - v_{EP}|\right) \quad (98)$$

Once we show that $v = v_{EP} + \mathcal{O}\left(n^{-2}\right)$, eq. (98) will give us indeed that $\mu = \mu_{EP} + \mathcal{O}\left(n^{-2}\right)$.

Let's now show that $v \approx v_{EP}$. We start from the first order expansion of m_2^{-1} from our extension of the Brascamp-Lieb theorem (eq. (48)). For $p(x)$, this gives us:

$$\left|v^{-1} - \phi_p^{''}(\mu)\right| \leq \frac{n^2K_3^2}{n^2\beta_m^2} + \frac{nK_4}{2n\beta_m}$$
$$\leq \frac{K_3^2}{\beta_m^2} + \frac{K_4}{2\beta_m} \quad (99)$$

Again the corresponding relationship for the hybrids is not exactly what we want it to be:

$$\forall i \quad \left|v_{EP}^{-1} - \phi_i^{''}(\mu_{EP}) - \beta_{-i}\right| \leq \frac{K_3^2}{n^2\beta_m^2} + \frac{K_4}{2n\beta_m}$$
$$\leq n^{-2}\frac{K_3^2}{\beta_m^2} + n^{-1}\frac{K_4}{2\beta_m} \quad (100)$$

But again, we sum all those relationships:

$$\left|nv_{EP}^{-1} - \phi_p^{''}(\mu_{EP}) - (n-1)\beta\right| \leq n^{-1}\frac{K_3^2}{\beta_m^2} + \frac{K_4}{2\beta_m} \tag{101}$$

which further simplifies, because $\beta = v_{EP}^{-1}$, into:

$$\left|v_{EP}^{-1} - \phi_p^{''}(\mu_{EP})\right| \leq n^{-1}\frac{K_3^2}{\beta_m^2} + \frac{K_4}{2\beta_m} \tag{102}$$

Again, we find that the pairs (μ, v) and (μ_{EP}, v_{EP}) obey very similar relationships: eqs. (99) and (102). We have:

$$\left|\phi_p^{''}(\mu) - \phi_p^{''}(\mu_{EP})\right| \leq K_3\left|\mu - \mu_{EP}\right| \tag{103}$$

and this gives us that $v^{-1} \approx v_{EP}^{-1}$:

$$\begin{aligned}
\left|v^{-1} - v_{EP}^{-1}\right| &\leq K_3\left|\mu - \mu_{EP}\right| + \left(1 + n^{-1}\right)\frac{K_3^2}{\beta_m^2} + 2\frac{K_4}{2\beta_m} \\
\left|v^{-1} - v_{EP}^{-1}\right| &\leq \mathcal{O}\left(1\right) + \mathcal{O}\left(n\left|\mu - \mu_{EP}\right|\right)
\end{aligned} \tag{104}$$

Our final equations for the size of $|\mu - \mu_{EP}|$ and $\left|v^{-1} - v_{EP}^{-1}\right|$ seem to be caught in a loop: you need to know how good one approximation is in order to know how good the second will be and so on. This is not at all the case and it is very easy to cut this loop.

The easiest way is to remark that both μ and μ_{EP} are $\mathcal{O}\left(n^{-1}\right)$ away from the mode of p and so they must be $\mathcal{O}\left(n^{-1}\right)$ from one another (see main text, section 2.2). This gives $v^{-1} = v_{EP}^{-1} + \mathcal{O}\left(1\right)$ (from eq. (104).

Then, we remark that both v^{-1} and v_{EP}^{-1} are order n. The error for $|v - v_{EP}|$ is then of order n^{-2} and we have that $\mu = \mu_{EP} + \mathcal{O}\left(n^{-2}\right)$, from eq. (94). This concludes the first part of our proof.

Let's now look at the fourth moment of the target m_4^p. We will show that is matched to by the fourth moment m_4^i of any hybrid and by the fourth moment of the Gaussian approximation of $p(x)$: $3v_{EP}^2$.

From our Brascamp-Lieb extension, the first order approximation of m_4^p is:

$$\left|\phi^{''}(\mu)m_4^p - 3v\right| \leq n^{-2}\left(\frac{19}{2}\frac{K_3^2}{\beta_m^4} + \frac{5}{2}\frac{K_4}{\beta_m^3}\right) \tag{105}$$

From which, intuitively: $m_4^p \approx 3v\left(\phi_p^{''}(\mu)\right)^{-1} \approx 3v^2 \approx 3v_{EP}^2$

Let's now formalize this intuition by bounding explicitely each error term:

$$\begin{aligned}
3v^2 - 3v_{EP}^2 &= 6\left(v - v_{EP}\right)\frac{\left(v + v_{EP}\right)}{2} && (106)\\
\left|3v^2 - 3v_{EP}^2\right| &\leq 6\left|v - v_{EP}\right|\frac{1}{2n\beta_m} && (107)\\
3v\left(\phi_p^{''}(\mu)\right)^{-1} - 3v^2 &= 3v\left[\left(\phi_p^{''}(\mu)\right)^{-1} - v\right] && (108)\\
\left|3v\left(\phi_p^{''}(\mu)\right)^{-1} - 3v^2\right| &\leq \left|\left(\phi_p^{''}(\mu)\right)^{-1} - v\right|\frac{3}{n\beta_m} && (109)
\end{aligned}$$

Which we can bound using preceding relationships (eq. (104) and eq. (49)), and which gives us the final bound:

$$\begin{aligned}
\left|m_4^p - 3v_{EP}^2\right| &\leq n^{-3}\left(\frac{19}{2}\frac{K_3^2}{\beta_m^5} + \frac{5}{2}\frac{K_4}{\beta_m^4}\right) + \frac{6}{n\beta_m}\left|v - v_{EP}\right| + \frac{3}{n\beta_m}\frac{1}{n^2\beta_m^2}\left[\frac{2K_3^2}{\beta_m^2} + \frac{K_4}{2\beta_m}\right] \\
&\leq \mathcal{O}\left(n^{-3}\right)
\end{aligned} \tag{110}$$

Let's note that this final approximation isn't any better of any worse, in terms of orders of magnitude, than the original approximation $m_4^p \approx 3v\left(\phi_p^{''}(\mu)\right)^{-1}$.

Another approximation that is of similar quality, in terms of orders of magnitude, is for any hybrid i: $m_4^p \approx m_4^i$. Indeed, from 51 (Brascamp-Lieb extension: m_4 first order approximation), we have that:

$$\left|\left[\phi_i^{''}(\mu_{EP}) + \beta_{-i}\right] m_4^i - 3v_{EP}\right| \leq n^{-4}\frac{19}{2}\frac{K_3^2}{\beta_m^4} + n^{-3}\frac{5}{2}\frac{K_4}{\beta_m^3} \tag{111}$$

and see that m_4^i would obey a similar relationship to m_4^p (eq. (105)) if $\beta_{-i} \approx \sum_{j\neq i}\phi_j^{''}(\mu_{EP})$. That happens to be the case because we also have:

$$\left|v_{EP}^{-1} - \left[\phi_i^{''}(\mu_{EP}) + \beta_{-i}\right]\right| \leq n^{-2}\frac{2K_3^2}{\beta_m^2} + n^{-1}\frac{K_4}{2\beta_m} \tag{112}$$

Thus, $\phi_i^{''}(\mu_{EP}) + \beta_{-i}$ is approximately constant (in i), and approximately equal to v_{EP}^{-1}, which is an important result in its own right. If we combine eqs. (111) and (112), we thus have:

$$\begin{aligned} m_4^i &= 3v_{EP}\left[\phi_i^{''}(\mu_{EP}) + \beta_{-i}\right]^{-1} + \mathcal{O}\left(n^{-4}\right) \\ &= 3v_{EP}^2 + \mathcal{O}\left(n^{-4}\right) \\ &= m_4^p + \mathcal{O}\left(n^{-3}\right) \end{aligned} \tag{113}$$

which concludes our proof that all fourth moments of the hybrids and q and p are approximately equal. Note that an absolute error of order n^{-3} translates into a relative error of order n^{-1}.

Let's now show how to approximate the third moment of the target m_3^p from the third moments of the hybrids m_3^i. We start for the first-order approximation of m_3^p (Brascamp-Lieb extension, eq. (50)):

$$\left|\phi_p^{''}(\mu)m_3^p + \left(\phi^{'}(\mu)v + \frac{\phi^{(3)}(\mu)}{2}m_4^p\right)\right| \leq n^{-2}\left(\frac{17}{6}\frac{K_3K_4}{\beta_m^4} + \frac{5}{8}\frac{K_5}{\beta_m^3}\right) \tag{114}$$

$$m_3^p \approx -\left(\phi_p^{''}(\mu)\right)^{-1}\left(\phi^{'}(\mu)v + \frac{\phi^{(3)}(\mu)}{2}m_4^p\right) \tag{115}$$

For the hybrids, we have:

$$\forall i \left|\left(\phi_i^{''}(\mu_{EP}) + \beta_{-i}\right)m_3^i + \left(\left(\phi_i^{'}(\mu_{EP}) + \beta_{-i}\mu_{EP} - r_{-i}\right)v_{EP} + \frac{\phi_i^{(3)}(\mu_{EP})}{2}m_4^i\right)\right| \leq n^{-4}\frac{17}{6}\frac{K_3K_4}{\beta_m^4} + n^{-3}\frac{5}{8}\frac{K_5}{\beta_m^3} \tag{116}$$

We will perform the following steps:

$$m_3^i \approx \left(\phi_i^{''}(\mu_{EP}) + \beta_{-i}\right)^{-1}\left(\left(\phi_i^{'}(\mu_{EP}) + \beta_{-i}\mu_{EP} - r_{-i}\right)v_{EP} + \frac{\phi_i^{(3)}(\mu_{EP})}{2}m_4^i\right) \tag{117}$$

$$\approx v_{EP}\left(\left(\phi_i^{'}(\mu_{EP}) + \beta_{-i}\mu_{EP} - r_{-i}\right)v_{EP} + \frac{\phi_i^{(3)}(\mu_{EP})}{2}m_4^p\right) \tag{118}$$

From which:

$$\sum_i m_3^i \approx v_{EP}\sum\left(\left(\phi_i^{'}(\mu_{EP}) + \beta_{-i}\mu_{EP} - r_{-i}\right)v_{EP} + \frac{\phi_i^{(3)}(\mu_{EP})}{2}m_4^p\right) \tag{119}$$

$$\approx v_{EP}\left(\left(\phi_p^{'}(\mu_{EP}) + 0\right)v_{EP} + \frac{\phi_p^{(3)}(\mu_{EP})}{2}m_4^p\right) \tag{120}$$

from which we see that m_3^p and $\sum_i m_3^i$ obey very similar relationships (eq. (114) and eq. (120)), and can conclude that they are close.

More formally, starting from eq. (116), let's replace $\phi_i^{''}(\mu_{EP}) + \beta_{-i}$ with $\phi_p^{''}(\mu_{EP})$:

$$\begin{aligned} m_3^i &= \left(\phi_i^{''}(\mu_{EP}) + \beta_{-i}\right)^{-1} \left(\left(\phi_i^{'}(\mu_{EP}) + \beta_{-i}\mu_{EP} - r_{-i}\right) v_{EP} + \frac{\phi_i^{(3)}(\mu_{EP})}{2} m_4^i \right) + \mathcal{O}\left(n^{-4}\right) \\ &= v_{EP} \left(\left(\phi_i^{'}(\mu_{EP}) + \beta_{-i}\mu_{EP} - r_{-i}\right) v_{EP} + \frac{\phi_i^{(3)}(\mu_{EP})}{2} m_4^i \right) + \mathcal{O}\left(n^{-2}n^{-2}\right) + \mathcal{O}\left(n^{-4}\right) \end{aligned} \tag{121}$$

Now, we replace m_4^i with m_4^p. Since, $m_4^i = m_4^p + \mathcal{O}\left(n^{-3}\right)$, we have:

$$m_3^i = v_{EP} \left(\left(\phi_i^{'}(\mu_{EP}) + \beta_{-i}\mu_{EP} - r_{-i}\right) v_{EP} + \frac{\phi_i^{(3)}(\mu_{EP})}{2} m_4^p \right) + \mathcal{O}\left(n^{-4}\right) \tag{122}$$

which we finally sum for i: the $\beta_{-i}\mu_{EP} - r_{-i}$ sum to 0, leaving:

$$\sum_i m_3^i = v_{EP} \left(\phi_p^{'}(\mu_{EP}) v_{EP} + \frac{\phi_p^{(3)}(\mu_{EP})}{2} m_4^p \right) + \mathcal{O}\left(n^{-3}\right) \tag{123}$$

Because, $\mu = \mu_{EP} + \mathcal{O}\left(n^{-2}\right)$ and $v = \left(\phi_p^{''}(\mu)\right)^{-1} + \mathcal{O}\left(n^{-2}\right) = v_{EP} + \mathcal{O}\left(n^{-2}\right)$, $\sum_i m_3^i$ and m_3^p have identical first order expansions (which is of order n^{-2}). More precisely:

$$\begin{aligned} \phi_p^{'}(\mu_{EP}) v_{EP} + \frac{\phi_p^{(3)}(\mu_{EP})}{2} m_4^p &= \phi_p^{'}(\mu_{EP}) v + \frac{\phi_p^{(3)}(\mu_{EP})}{2} m_4^p + \mathcal{O}\left(n^{-2}\right) &(124) \\ &= \phi^{'}(\mu) v + \frac{\phi^{(3)}(\mu)}{2} m_4^p + \mathcal{O}\left(n^{-2}\right) + \mathcal{O}\left(|\mu - \mu_{EP}|\right) &(125) \end{aligned}$$

because: $\left|\phi^{'}(\mu) - \phi^{'}(\mu_{EP})\right| \leq n\beta_m |\mu - \mu_{EP}|$ and, similarly, $\phi^{(3)}(\mu) - \phi^{(3)}(\mu_{EP}) = \mathcal{O}\left(n|\mu - \mu_{EP}|\right)$. And:

$$\begin{aligned} (v - v_{EP}) \left(\phi_p^{'}(\mu_{EP}) v_{EP} + \frac{\phi_p^{(3)}(\mu_{EP})}{2} m_4^p \right) &= \mathcal{O}\left(|v - v_{EP}|\right) \left(\mathcal{O}(1)\mathcal{O}\left(n^{-1}\right) + \mathcal{O}(n)\mathcal{O}\left(n^{-2}\right) \right) \\ &= \mathcal{O}\left(n^{-3}\right) \end{aligned} \tag{126}$$

Which gives us the final expression:

$$m_3 = \sum_i m_3^i + \mathcal{O}\left(n^{-3}\right) \tag{127}$$

which concludes our proofs on the quality of the EP approximation.

In the main text, we have also used the following relationship, detailing the second order expansion of v_{EP}^{-1}:

$$v_{EP}^{-1} = \phi_p^{''}(\mu_{EP}) + \sum_i \left[\phi_i^{(3)}(\mu_{EP}) \frac{m_3^i}{2v_{EP}} \right] + \phi_p^{(4)}(\mu_{EP}) \frac{3v_{EP}^2}{3! v_{EP}} + \mathcal{O}\left(n^{-1}\right) \tag{128}$$

For the inquisitive reader, this is obtained by starting from our Brascamp-Lieb extension, eq. (55), applied to all hybrids. Then proceeding to approximate $m_4^i \approx 3v_{EP}^2$ and summing. □